University Collaborative Research Program

The ISCR supported seven University Collaborative Research Program (UCRP) projects during FY 2004 at the University of California campuses. These projects support graduate students working on doctoral thesis research, and the faculty principal investigators and students work closely with an LLNL collaborator. The projects in FY 2004 spanned four different UC campuses, as well as the departments of civil engineering, computer science, mechanical engineering, and physics. They ranged from understanding and optimizing the performance of advanced architecture computers to improving engineering simulation models; from the development of tracking and recognizing objects to basic computational science at the frontier.

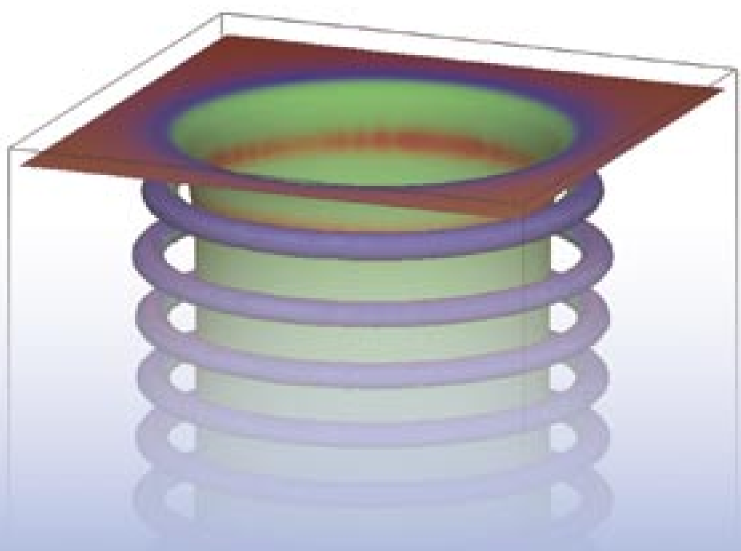

Data-Driven Execution of Communication-Tolerant Algorithms

Principal Investigator **Scott B. Baden,** University of California, San Diego

Collaborator **Daniel Quinlan,** CASC

Communication-tolerant algorithms are expected to play an important role in achieving scalability on large-scale platforms with many thousands of processors. However, communication-tolerant application development is time-consuming and prone to error, even for the expert programmer. A run-time substrate that supports non-bulk, synchronous, data-driven execution is under investigation. Because it treats communication and computation as coupled activities, rather than as distinct phases of execution, the data-driven execution model naturally supports irregular or unpredictable communication delays that are expected on large-scale platforms, such as BlueGene/L.

The data-driven substrate programmer expresses an application in terms of partially ordered communication and computation operations, as constrained by the data dependences inherent to the application. The substrate invokes an externally specified scheduler to determine the precise task-execution ordering. The user may further constrain the scheduling by means of *performance metadata* decorating the graph. The user expresses partially ordered operations in terms of a directed graph called a *task graph*. Vertices in the task graph correspond to tasks and edges to data dependences between the tasks.

The substrate is implemented as a C++ library and runs as a background thread (i.e., a proxy) to manage dependence information and task scheduling. This background thread, called the Mover–Dispatcher, routes data from completed tasks to dependent tasks and determines when such tasks are enabled for execution.

When several possible tasks are ready for execution, their exact order is determined by the scheduler in conjunction with performance metadata. We built the communication proxy part of the Mover–Dispatcher, which supports non-blocking, asynchronous communication by means of multithreading. We are currently testing this capability on a latency-tolerant variant of a Gauss Seidel red-black Poisson solver in three dimensions using a Beowulf cluster in our research lab.

The interpretation of performance metadata is up to the application. For example, it might express priorities or affinities. We have tested out this capability on blocked LU running on a single processor. Our goal was to reorder computations in order to improve memory locality. The blocked LU application demonstrated the importance of an economical run-time representation for the task graph structure. The overhead of dynamically reordering task execution is 25% of the overall running time and an improved run-time representation for tasks graph representation is under investigation. We are exploring two optimizations. The first takes advantage of graph sparsity—most nodes depend on only a few other nodes. The second takes advantage of temporal locality. Rather than instantiate an entire task graph at initialization time, another approach is to create task graph nodes on demand and to recycle disused nodes. This strategy effectively throttles task generation and avoids tying up memory with task information that won't be needed for a long time.

Publications

S. B. Baden, "Moving forward in large scale computation" in J Dongarra, K. Madsen, J. Wasniewski (Eds.), 7th International Conference on Applied Parallel Computing (PARA '04), Lyngby Denmark, June 2004, *Lecture Notes in Computer Science,* Springer, 2004.

S. B. Baden, "Masking Latency with Data Driven Program Variants" in J Dongarra, K. Madsen, J. Wasniewski (Eds.), 7th International Conference on Applied Parallel Computing (PARA '04), Lyngby Denmark, June 2004, *Lecture Notes in Computer Science,* Springer, 2004.

Feature-Based Approaches for Long-Range Motion Segmentation and Object Tracking

Principal Investigator **Serge Belongie,** University of California, Davis

Collaborator **Chandrika Kamath,** CASC

Image segmentation is the problem of partitioning the pixels in an image into a relatively small number of regions that correspond to objects or parts of objects. It is one of the hardest (and oldest) open problems in computer vision, and it plays an important role in the process of object detection and recognition. As challenging and computationally intensive as image segmentation is, it also happens to be a problem that the human visual system solves effortlessly. The goal of this project is to develop methods for image and video segmentation with an emphasis on motion-based processing.

Our proposed work divides into two main areas: "what went where" and "who went where." The first area, on which we have already begun working, addresses the problem of motion segmentation for image sequences with large inter-frame displacement, e.g., more than 10% of the image width. The second area, for which the principal application area will be people tracking, deals with the problem of detecting and tracking moving objects that exhibit temporally periodic variation in appearance.

We first describe our solution to the "what went where" problem. The algorithm operates in two stages, starting with robust estimation of the underlying motion fields and concluding with dense assignment of pixels to motion fields. The first stage of this process is the first dense-motion segmentation method to operationalize the layer-based formulation for multiple discrete motions. We detect the motion layers using a variant of RANSAC (Random Sample Consensus) on detected interest points using a planar projective (2D homography) motion model. We perform the layer assignment using a fast, graph-cut-based MRF (Markov Random Field) formulation, which enforces spatially piecewise smooth-pixel assignments.

Illustration of video object deletion. Original frames (top). Segmented layer corresponding to hand motion (middle). Reconstruction without the hand layer using the recovered motion of the keyboard (bottom). Note that no additional frames beyond the three shown were used as input.

The case of moving objects that do not fit the planar projective model is addressed in the second stage. A poor fit to the planar projective model can arise from two main causes: non-planarity and non-rigidity. For this purpose, we refine the fit via iterative regression using a regularized thin-plate spline (TPS).

Our approach to the "who went where" problem builds on our approach to the first problem by interfacing it in a novel way with the classical tools of multiview geometry. In particular, we apply the knowledge that sets of frames depicting a periodically moving object in a given shape but with varying poses will (approximately) satisfy the multiview constraints for a rigid object. After automatically detecting the period of a given object in the field of view, we apply the "what went where" framework to segment the object from the background. The segmented foreground object across the period-separated frames is then treated as a conventional stereo- or multiview-reconstruction problem. Using this framework, we demonstrated a 3D reconstruction of sparse interest points on a video sequence of a pedestrian and showed preliminary results computing dense disparity using a graph-cut-based stereo correspondence engine.

Publications

Sameer Agarwal, Satya Mallick, David Kriegman and Serge Belongie, "On Refractive Optical Flow," *European Conference on Computer Vision*, Prague, Czech Republic, pp. 483–494, **vol. 2**, 2004.

Josh Wills and Serge Belongie, "A feature based method for determining dense long range correspondences," *European Conference on Computer Vision*, Prague, Czech Republic, pp. 170–182, **vol. 3**, 2004.

Serge Belongie and Josh Wills, "Structure from Periodic Motion," *International Workshop on Spatial Coherence for Visual Motion Analysis*, Prague, Czech Republic, 2004.

DNS and Modeling of Dispersion of Solid or Liquid Particles in Turbulent Flows

Principal Investigator **Said Elghobashi,** University of California, Irvine

Collaborator **Robert Lee,** Atmospheric Science Division

The release and subsequent transport of toxic materials into the atmosphere or inside and around buildings are of major concern to national security. These toxic materials can be in the form of gas, liquid droplets or solid particles and often undergo chemical reactions during the transport process. An important distinguishing feature of the transport of toxic materials into the atmosphere or inside and around buildings is the wide spectra of length scales and time scales involved. Length scales range from microns (particle size) to kilometers (turbulent eddy size) and the corresponding time scales range from microseconds to days. The approach of the proposed research is to employ direct numerical simulation (DNS) for prototypical particle-laden turbulent flows (e.g., flow over a backward-facing step) to evaluate the turbulent *correlations directly* and provide accurate models for them that can be introduced in LES (Large Eddy Simulation) or RANS (Reynolds-Averaged Navier–Stokes equations) codes.

The proposed work encompasses the numerical solution of the three-dimensional, time-dependent Navier–Stokes and continuity equations of the turbulent flow in addition to solving the equations of motion of the dispersed particles. The time-averaged correlations needed for RANS equations will be evaluated from their instantaneous values in DNS via ensemble and time averaging, and then models will be developed to relate these correlations to the dependent variables used in RANS equations.

Our work on the project during the past six months concentrated on the mathematical development of the necessary boundary conditions for the DNS code for the backward-facing step and performing short runs on a small parallel computer. We also spent considerable effort in converting our DNS parallel code that was running on a Cray-T3E so that it can run on a MCR computer (LLNL). A new machine-independent Fast Fourier Transform (FFT) was written and incorporated into the code. We are starting the DNS of the flow over the backward-facing step.

Numerical Study of Coexisting Superconductivity and Ferromagnetism: Applications to Real Materials

Principal Investigator **Warren E. Pickett,** University of California, Davis

Collaborator **Francois Gygi,** CASC

The discovery of ferromagnetic metals that become superconducting at lower temperatures (UGe_2, $ZrZn_2$, URhGe) early in this young century provided one of the big surprises in materials physics in recent years. For more than four decades, such coexistence was believed to be all but impossible. This coexistence of two competing phases, each a macroscopic manifestation of the quantum behavior of electrons, fully qualifies this as a new state of matter. The goal of this project is to include material-specific characteristics of the metallic states in the underlying formalism and to perform numerical studies to illuminate the microscopic driving mechanisms.

Our approach consists of three prongs:

(1) Reformulating the theory at its more basic levels to gain insight

(2) Adapting the formalism, developing algorithms and making the resulting codes applicable to real materials, such as those mentioned above, as opposed to the earlier study of model systems

(3) Computing solutions to the superconducting gap equations to map out phase diagrams.

Our work has moved to simulating the novel type of new coexisting ferromagnetic-superconducting state called FFLO, after its discoverers Fulde, Farrell, Larkin and Ovchinnikov. This state, which can occur in very weak ferromagnets or for paramagnets near the critical value of magnetic field, consists of the coalescence of superconducting pairs of electrons with a non-vanishing pair momentum. The resulting state and its order parameter are inhomogeneous in space. Our primary results are the identification of phase boundaries that specify the phase diagram.

The results of this year's research are:

(1) Identification of a characteristic velocity spectral function that is a fundamental part of the integral of the integral equation

(2) Development and implementation of an algorithm to evaluate this function

(3) Showing how its spectrum affects the phase diagram and obtain predictions for a specific candidate, $ZrZn_2$.

During the past year, two new materials have been found to display some of the characteristics of the FFLO phase. These two systems are the peculiar rare-earth-based superconductor $CeCoIn_5$ and one of the first "heavy fermion" superconductors, UBe_{13}. If confirmed to be FFLO, these will add to the single system (the two-dimensional organic metal k-$(BEDT\text{-}TTF)_2Cu(NCS)_2$) that is relatively well established to enter an FFLO phase when a magnetic field is applied. One strong focus of the coming year will be to apply our new methods to these novel superconductors and help assess whether an FFLO phase really does arise. The computational facilities of ISCR are anticipated to be important for the success of the latter stages of this project.

Publications

A. B. Kyker, W. E. Pickett, and F. Gygi, "Fermiology and Fulde-Farrell-Larkin-Ovchinnikov Phase Formation," *Physical Review B* (submitted), 2004.

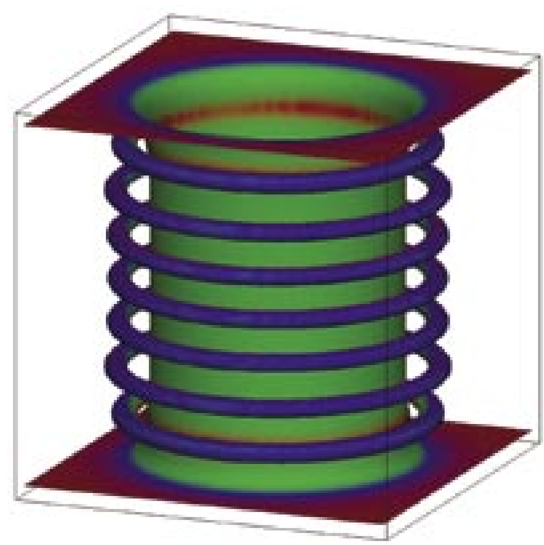

A current-carrying wire penetrates a slab of superconductor where it is electrically insulated, producing a radial magnetic field. Such a field, penetrating a type-II superconductor, produces regions (flux lines) of depressed superconducting strength, each containing one quantum of magnetic flux. For this geometry, the flux lines align regularly in rings, a generalization of the Abrikosov vortex lattice state in superconducting films. The color here gives the rate of change of the superconducting order parameter on the various isosurfaces (red is slow, blue is rapid).

Lagrangian Simulation of Penetration and Other Extreme-Deformation Events: Moving Beyond Meshless Methods

Principal Investigator **Mark Rashid,** University of California, Davis

Collaborator **Mike Puso,** Engineering

Solid mechanics problems involving extreme-deformation events, such as projectile/target interaction and fragmentation of cased explosives, present a considerable challenge to the conventional Lagrangian finite element method (FEM). Accordingly, beginning in the mid-1990s, various so-called meshless approximation methods have been proposed as possible alternative approaches. This research project involves the development of a variational approximation method that attempts to combine the best features of both the conventional FEM and meshless methods.

The technical approach involves a synergy of two distinct and innovative elements. In the first of these, the solid-mechanics problem is discretized using a flexible geometric subdivision scheme in which "elements" take the form not of hexahedra or tetrahedra but of arbitrary polyhedra. This flexibility vastly simplifies the task of mesh generation and thereby facilitates periodic remeshing, as is often required by extreme mesh distortion. The second element of the project involves a novel method for transferring information relating to the current material state from the old, distorted mesh to the new mesh following a remeshing cycle. This *remapping* step is required before the solution can be continued with the new mesh and can be a significant source of errors.

VETFEM is a finite-element-like Galerkin approximation method in which the basis functions and numerical quadrature are facilitated by a polyhedral partition of the material domain. In contrast to the conventional FEM, wherein all elements must conform to strict geometric and topological requirements, VET elements can take arbitrary polyhedral forms. The VET element formulation has been successfully extended to 3D, coded, and integrated into an existing research code. This code offers flexible and expandable capabilities, including explicit dynamics, implicit dynamics, quasistatics, strongly-objective (Rashid, 1993) finite-deformation one.

As with the VET element formulation, the extension of the existing 2D state remapping methodology to 3D turned out to be challenging. The fundamental algorithmic problem associated with the variational remapping method is that of partitioning the volume of a "new" tributary region among a tiling of "old" tributary regions with which it overlaps. The key innovation involves a highly efficient and robust optimization procedure that closely approximates the required volume partition without resorting to any cumbersome geometric calculations. This has been accomplished and the result is a highly efficient and robust piece of code.

A preliminary result is illustrated in figure, which shows a twisted toroidal mesh consisting of eight-node hex elements. A fairly complex function was defined on this mesh, and then it was rotated about its axis through 90 degrees in a series of steps. After each increment of rotation, the integration-point values of the function were remapped from the "old" mesh to the new one. Because the function itself possessed a periodicity of 90 degrees, the final, succession can be directly compared to the original ones. The error for different numbers of rotation increments is shown in the table. The performance of the new remapping method is being studied with this and other types of analyses.

Twisted toroidal mesh is rotated 90 degrees in a series of steps, with remapping occurring after each increment of rotation. The error is shown for a range of the number of steps.

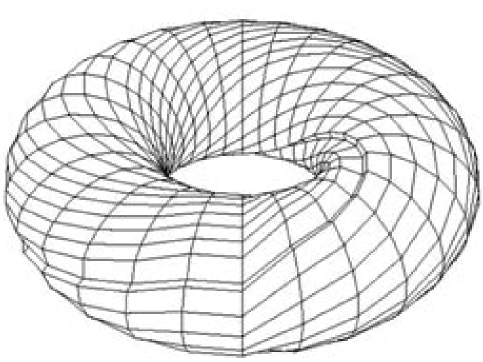

No. of rotation increments	RMS error	Max error
1	0.013404	0.027977
2	0.0132459	0.0441313
3	0.0148609	0.0547436
4	0.0174637	0.0616309
5	0.0193817	0.0616309
6	0.0212833	0.07386956
7	0.0228581	0.0830038
8	0.0244043	0.087037
9	0.02591708	0.1009022
10	0.027899	0.166627

Publications

M.M. Rashid and M. Selimotic "A Three-Dimensional Finite Element Method With Arbitrary Polyhedral Elements," submitted, 2004

Memory Access Pattern Signatures and Certificates of Relevance for Benchmarks

Principal Investigator **Allan Snavely,** San Diego Supercomputer Center and University of California, San Diego

Collaborator **Bronis de Supinski,** CASC

This award has supported Michael O. McCracken in his fundamental research investigations into performance modeling and prediction as part of a collaboration involving the University of California, San Diego's Computer Science Department (where McCracken is a student and Allan Snavely an adjunct assistant professor), the San Diego Supercomputer Center's (SDSC's) Performance Modeling and Characterization (PMaC) laboratory led by Snavely, and LLNL's Bronis de Supinski. Michael has thus been able to participate in research relevant to both organizations and also to PERC (the Performance Evaluation Research Center), a DOE Office of Science ISIC.

So far our research has investigated the performance implications of memory access patterns and useful definitions of "signature distance" between the memory access patterns of different basic blocks from the same or different programs. The goal is to improve the accuracy and speed of the Convolution method for performance prediction. We have enhanced the functionality of the MetaSim tool for gathering memory access pattern signatures and have made this tool platform independent. We have been investigating what kinds of memory access patterns exist "in nature" and exploring the performance implications of memory access patterns. We have developed a nomenclature and symbolic representation of memory access patterns leveraging previous work by Nick Mitchell.

We have been examining definitions of "signature distance" between basic blocks—our expectation is that basic blocks with similar memory access patterns will perform similarly on a given machine. In work defining a meaningful metric for "signature distance," he established orthogonal properties of loops including memory footprint, memory access pattern, type and intensity of floating-point, and ILP (instruction level parallelism) operations that could provide "certificates of relevance" for benchmarks. This work has demonstrated that reasonably accurate estimates of an application's performance can be derived from its MetaSim profile and mapping its basic blocks to a similar set of benchmark basic blocks with an established performance profile. McCracken has published a paper showing how these performance model predictions can guide dynamic algorithm selection.

Currently McCracken is developing a study of the HPC Challenge Benchmarks that will position these, and several strategic DOE applications, in dimensions of spatial and temporal locality. This work was presented at the HPCS meeting at SC04 in Pittsburgh. Further direct collaboration with LLNL is resulting in the means to acquire the application profiles via static analysis using the ROSE compiler infrastructure. This approach can result in significant time-savings and is leveraging de Supinski's expertise in compiler technology. The idea is that much of the information currently acquired via tracing, such as operation types and counts, communication and memory access patterns, can be determined or at least reasonably estimated more rapidly via static analysis. Tracing will then be used simply to confirm compile-time information via sampling and to fill in information (such as loop bounds) that may in some cases be unknown at compile time. The goal is to speed up code profiling by an order of magnitude.

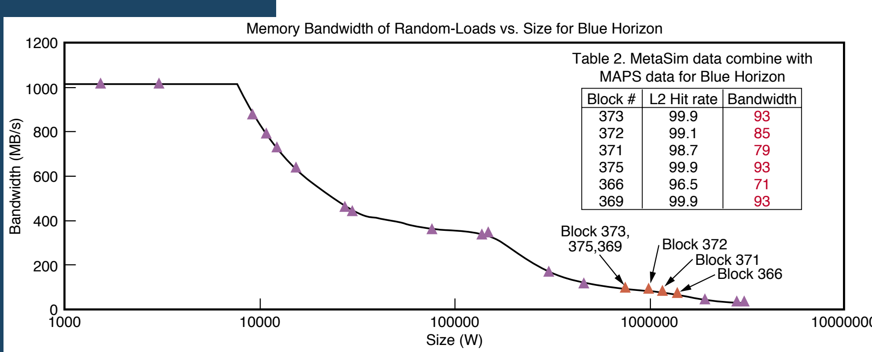

Horizon Machine Profile with MetaSim data.

Publications

M.O. McCracken, A. Snavely, A. Malony, "Performance Modeling for Dynamic Algorithm Selection," *ICCS Workshop on Performance Modeling and Analysis (PMA03)*, June 2003, Melbourne, Australia.

Visual Tracking and Recognition for Biometrics and Interactive Visualization

Principal Investigator **Matthew Turk,** University of California, Santa Barbara

Collaborator **Lenny Tsap**, Electronics Engineering Technologies

At the Four Eyes Lab in the Computer Science Department at the University of California, Santa Barbara, we have pursued research on visual tracking and recognition for biometrics and interactive visualization with notable progress and success. Computer vision is a promising and powerful sensing modality that can be used to unobtrusively track, model, and classify human appearance and behavior.

Our primary goals for the year were to make progress in

1. Face tracking and facial expression analysis
2. Hand detection, tracking, and gesture recognition
3. Extracting and using reliable depth edges in images and video.

These areas contribute to the overall goals of the project and give us a solid start in addressing the technical needs of biometrics and visualization problems. In face-tracking and facial-expression analysis, we extended our wavelet-based method to enable real-time tracking, and we used embedding techniques to represent high-dimensional face information in low-dimensional manifolds to investigate the representation and recognition of dynamic facial expressions [1, 2]. We implemented and tested real-time hand tracking and recognition and successfully applied it to mobile computing applications [3,4], winning a workshop Best Paper award. We also developed a technique for reliably extracting depth edges using a digital camera with multiple flashes and developed methods to produce non-photorealistic renderings of scenes and recognize "fingerspelling" letters (used in sign language) from images, which resulted in additional publications [5,6].

In the past decade, there has been an increasing interest in moving beyond the traditional applications of computer vision, such as robot navigation, object recognition, and industrial inspection, to using vision technology as an effective input modality in human-computer interaction (HCI) often pursued in the context of multimodal or perceptual interfaces. The general focus of these efforts is to integrate multiple perceptual modalities, such as computer vision, speech and sound processing, and haptic I/O, into the user interface. Such video-based sensing is passive and non-intrusive as it does not require contact with the user or any special-purpose devices.

The primary tasks of computer vision in these scenarios are to detect, model, recognize, and interpret various visual aspects of human behavior. If delivered reliably and robustly, such vision technology can support a range of functionality in

Hand tracking using the "Flocks of Features" method.

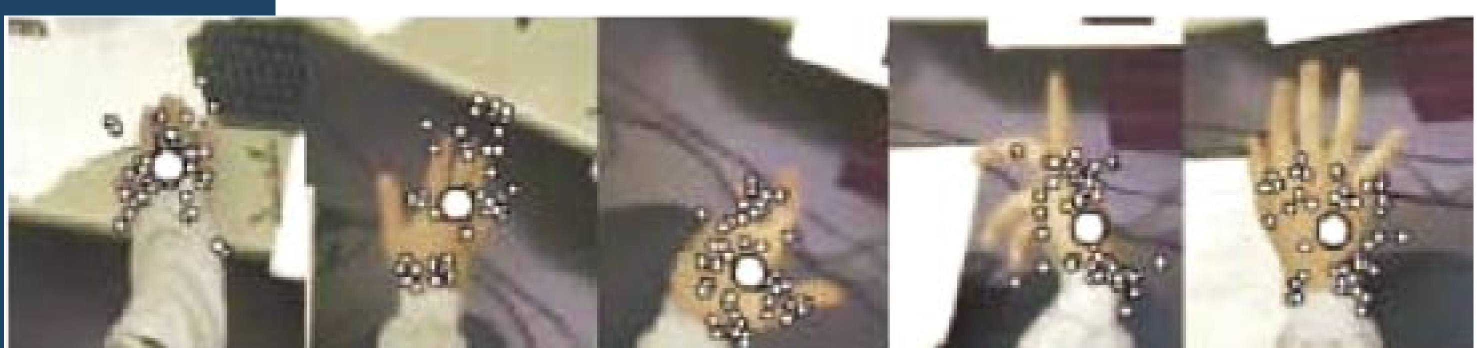

interactive systems by conveying relevant visual information about the user, such as identity, location, and movement, thus providing key contextual information. In order to fully support visual aspects of interaction, several tasks must be addressed, such as face detection and recognition, head and face tracking, facial expression analysis, eye-gaze tracking, body tracking, hand tracking, gait recognition, and recognition of postures, gestures, and overall activity.

Current biometric systems are based on processing images of a user's face, iris pattern, hand geometry, fingerprint, and other physical characteristics. As dynamic tracking and recognition technologies mature, behavior may well become an important aspect of biometrics – e.g., how people make facial expressions, how they move, their gaze and blinking patterns, and their postures. Leveraging these technologies will require a different approach to biometrics, with integration across modalities and across time (temporal integration) becoming central to the problem of verification.

Publications

1. C. Hu, Y. Chang, R. Feris, M. Turk, "Manifold based analysis of facial expression," *IEEE Workshop on Face Processing in Video*, Washington, D.C., June 2004.

2. Y. Chang, C. Hu, M. Turk, "Probabilistic expression analysis on manifolds," *International Conference on Computer Vision and Pattern Recognition*, Washington DC, June 2004.

3. M. Kölsch and M. Turk, "Fast 2D hand tracking with flocks of features and multi-cue integration," *IEEE Workshop on Real-Time Vision for Human-Computer Interaction*, Washington DC, USA, June 2004 (Best Paper Award).

4. M. Kölsch and M. Turk, "Analysis of rotational robustness of hand detection with rectangle features," *International Conference on Pattern Recognition*, Cambridge, U.K., August 2004.

5. R. Feris, M. Turk, R. Raskar, K. Tan and G. Ohashi, "Exploiting depth discontinuities for vision-based fingerspelling recognition," *IEEE Workshop on Real-Time Vision for Human-Computer Interaction*, Washington DC, USA, June 2004.

6. R. Raskar, K. Tan, R. Feris, J. Yu, M. Turk, "A non-photorealistic camera: depth edge detection and stylized rendering using multi-flash imaging," *ACM SIGGRAPH*, Los Angeles, August 2004.

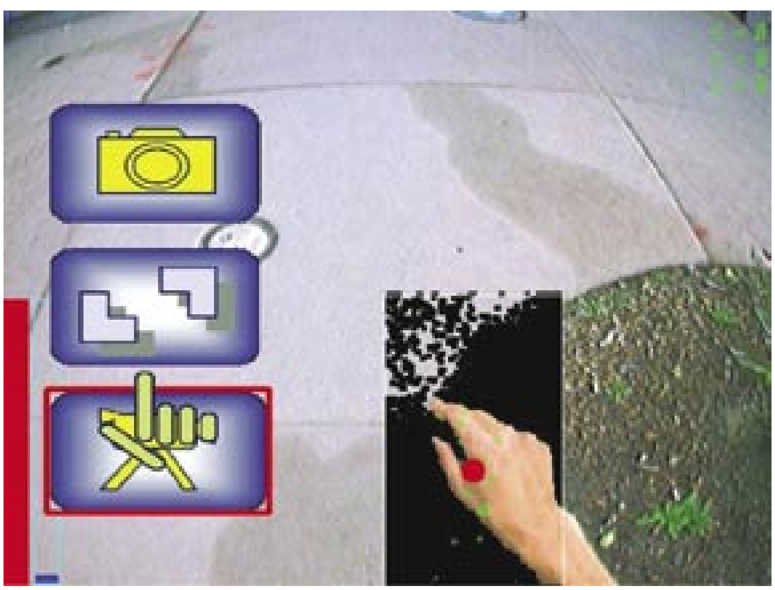

Figure 2. A user interface employing hand gestures.